

# Database Languages and their Compilers

Prof. Dr. Torsten Grust

Database Systems Research Group  
U Tübingen

Winter 2010



Copyright © 1995 United Feature Syndicate, Inc.  
Redistribution in whole or in part prohibited

# 0 Introduction to Query Compilation

## 0.1 Welcome!

You will have a hard time ...

...to find a more interesting and fun enterprise in Computer Science than the study of the innards of a **Database Query Processor**.

### Milestones:

- ▶ Dissect *all* phases of an SQL query processor;
- ▶ discuss the necessary theoretical toolbox;
- ▶ provide the details needed to get an actual implementation going.

In some sense this is a course on **Compiler Construction**, yet the input language and the generated code might appear strange:

### Sample Input (SQL):

```
SELECT COUNT(z), x, y
FROM S
GROUP BY x, y
```

### Sample Output (Algebraic Program):

```
1 Insert('S', InlineLoad(['x', 'y', 'z'],
2     '1|foo|1
3     1|foo|1
4     0|bar|0
5     ',
6     '|', '\n'));
7 Commit;
8 Aggr(MScan('S',
9     ['x', 'y', 'z']),
10    [x, y],
11    [cnt = count()]);
```

## 0.2 Administrativa

▶ **Lecture time slot** (OK everybody?):

- Wednesday, 10:15–11:45, in Sand 6/7 (kleiner Hörsaal)

▶ **Contact me:**

- Room B 318 (office hours: Thursday, 13:30–15:00)
- e-mail: `Torsten.Grust@uni-tuebingen.de`

▶ **Homepage for this course:**

- <http://www-db.informatik.uni-tuebingen.de/teaching/ws1011/dblang>

▶ **An oral exam** will conclude this course. Individual appointments.

## 0.3 Some Remarks on these Slides

This is (partly) new material I am preparing as the semester goes on; I will try to make all slides available prior to the lectures. (Please bear with me.)

### Notes

We will use non-slide material (software, blackboard) and references to such material will be marked like this. Check **your** notes.

### Attention

Boxes like this mark material of core interest and it might be advisable to double-check your understanding of the contents.

### If you have time ...

... you might be interested to follow references in sections like this to gain a deeper understanding of the subject (bedtime reading).

## Commercial Systems

Commercial DBMSs<sup>1</sup> sometimes come with interesting query processing technology. These boxes try to point that out.

## Caveat

Material in these sections might be tricky to understand (e.g., an algorithm that might not behave as you would guess on first sight)

## Woops!

...even more tricky stuff! Should not occur too often, though.

---

<sup>1</sup>If you want a copy of IBM DB2, feel free to ask me.

## 0.4 Relational Databases and SQL

Let us go through the following to align your and mine vocabulary and explain some fundamental issues.

- ▶ Relational database systems represent all user data in tabular structures.
- ▶ A specific **relation** (or **table**)  $S$  might be depicted like this:

$S$	$x$	$y$	$z$
	TRUE	'foo'	1
	TRUE	'foo'	1
	FALSE	'bar'	0

- $x$ ,  $y$ , and  $z$  are the **attributes** of relation  $S$ ;
- the possible values in the columns are restricted by the **attribute type**,  $x$  has type  $\mathbb{B}$  (boolean), we will often write  $x :: \mathbb{B}$ ; here, we also have  $y :: \mathbb{S}$  (string) and  $z :: \mathbb{N}$  (numeric);

$S$	$x$	$y$	$z$
	TRUE	'foo'	1
	TRUE	'foo'	1
	FALSE	'bar'	0

- ▶  $S$  contains 3 **tuples**, we say  $(\text{TRUE}, \text{'foo'}, 1) \in S$  or even  $(x : \text{TRUE}, y : \text{'foo'}, z : 1) \in S$  if we need to refer to attribute names; all tuples in  $S$  are of **tuple type**  $(x : \mathbb{B}, y : \mathbb{S}, z : \mathbb{N})$ ;
- ▶ the tuple type  $(x : \mathbb{B}, y : \mathbb{S}, z : \mathbb{N})$  is often called the **schema** of  $S$ ;
- ▶  $S$  may contain multiple identical tuples ( $S$  is a **multiset** or **bag** of tuples);
- ▶ we thus have  $S :: \text{bag}(x : \mathbb{B}, y : \mathbb{S}, z : \mathbb{N})$ , but we will also consider tables  $S :: \text{set } \alpha$  and  $S :: \text{list } \alpha$  with  $\alpha$  denoting any tuple type;
- ▶ in a database, we assume relation names to be unique; attribute names, however, may be re-used in different relations.



**SQL** (Structured Query Language) is the de-facto standard query interface offered by RDBMSs.

► SQL comes with language constructs to

■ selectively access only specific columns of a table:

SELECT $x, y$	---	<table border="1"><thead><tr><th><math>x</math></th><th><math>y</math></th></tr></thead><tbody><tr><td>TRUE</td><td>'foo'</td></tr><tr><td>TRUE</td><td>'foo'</td></tr><tr><td>FALSE</td><td>'bar'</td></tr></tbody></table>	$x$	$y$	TRUE	'foo'	TRUE	'foo'	FALSE	'bar'
$x$	$y$									
TRUE	'foo'									
TRUE	'foo'									
FALSE	'bar'									
FROM $S$										

■ evaluate filter predicates for every tuple:

SELECT *	---	<table border="1"><thead><tr><th><math>x</math></th><th><math>y</math></th><th><math>z</math></th></tr></thead><tbody><tr><td>FALSE</td><td>'bar'</td><td>0</td></tr></tbody></table>	$x$	$y$	$z$	FALSE	'bar'	0
$x$	$y$	$z$						
FALSE	'bar'	0						
FROM $S$								
WHERE $z \neq 1$								

■ aggregate table data:

```
SELECT SUM(z)
FROM S
```

 $\dashrightarrow$ 

1
2

■ join separate tables on given predicates:

```
SELECT s.x, s.z, t.z, t.y
FROM S AS s, S AS t
WHERE s.z < t.z
```

 $\dashrightarrow$ 

<i>x</i>	<i>z</i>	<i>z1</i>	<i>y</i>
FALSE	0	1	'foo'
FALSE	0	1	'foo'

■ group tuples with respect to given attributes and then aggregate over the resulting groups:

```
SELECT x, y, COUNT(z)
FROM S
GROUP BY x, y
```

 $\dashrightarrow$ 

<i>x</i>	<i>y</i>	3
TRUE	'foo'	2
FALSE	'bar'	1

- sort the tuples of a table by specified sort criteria:

<i>x</i>	<i>y</i>	<i>z</i>
FALSE	'bar'	0
TRUE	'foo'	1
TRUE	'foo'	1

- apply bag operations (UNION, EXCEPT, INTERSECT):

<i>x</i>	<i>y</i>	<i>z</i>
TRUE	'foo'	1
TRUE	'foo'	1
FALSE	'bar'	0
FALSE	'bar'	0

- remove duplicate tuples in a table:

```

SELECT DISTINCT *
FROM S
-->


| <i>x</i> | <i>y</i> | <i>z</i> |
|----------|----------|----------|
| TRUE     | 'foo'    | 1        |
| FALSE    | 'bar'    | 0        |


```

- evaluate universally or existentially quantified predicates:

```

SELECT *
FROM S
WHERE FORALL v IN VALUES 1, 2, 3 : v <> z
-->


| <i>x</i> | <i>y</i> | <i>z</i> |
|----------|----------|----------|
| FALSE    | 'bar'    | 0        |


```

#### 🏠 Standard SQL

Standard SQL requires you to write the above query as

```

SELECT *
FROM S
WHERE z <> ALL VALUES 1, 2, 3

```

- ▶ The result of applying a SQL query to database tables is a table again (SQL is a closed language).
- ▶ The SQL constructs `DISTINCT` and `ORDER BY` produce tables of type *set*  $\alpha$  and *list*  $\alpha$ , respectively. (N.B. the SQL processor might be completely *bag*-based internally.)

### 🕒 Relational DBMSs and SQL

You will find extensive coverage of relational database management systems (RDBMSs) and the SQL query language on the slides for the lecture “Datenbanksysteme I” (Summer 2010).

## Example 0.1

Suppose a cinema keeps track of seat reservations using the following database table Seats:

Seats	seat	reserved
	1	FALSE
	2	FALSE
	3	TRUE
	4	FALSE
	5	TRUE
	⋮	⋮

All seats in the cinema are represented in this relation, adjacent seats have adjacent seat numbers (we are ignoring the fact that seats  $n$  and  $n + 1$  might be in different rows).

 Use SQL to answer the following questions:

- ① How many seats are reserved already, how many are still free?
- ② List all groups of free seats of size 3.
- ③ List all groups of free seats of size  $n$  ( $n \geq 1$ ).

## 0.5 A Guided Tour through an SQL Query Processor

What follows provides a first insight into the workings of an SQL compiler and query processor. Consider this as a **roadmap for the rest of this course**. We will use the **example data and query** below and proceed step by step, producing a **list of tasks** we'll have to tackle as we go.

Mail	sender	subject	domain	DNS	domain	country
	'Alice'	'I love you'	' .nl'		' .de'	'Germany'
	'Bob'	'Re: You are fired'	' .uk'		' .nl'	'Netherlands'
	'Cathy'	'You are fired'	' .uk'		' .uk'	'United Kingdom'
	'Dave'	'I love you'	' .de'		⋮	⋮
	⋮	⋮	⋮			

```
SELECT Mail.sender, country
FROM Mail, DNS
WHERE Mail.domain = DNS.domain
AND subject = 'I love you'
```

## 0.5.1 Character Streams and Tokens

- ▶ The SQL processor does not see the query as a whole but consumes user input **character-by-character**:

```
S E L E C T   M a i l . s e n d e r ,   c o u n t r y '   F R O M   . . .   y o u ' ;
```

- Some characters do not contribute to the meaning of the query but rather define “word” boundaries:  `'`;
- it will be helpful for later stages to operate on words rather than characters; identify words and classify them: **tokens**.

S E L E C T	-->	<SELECT>
- 4 2	-->	<NUM, -42>
M a i l	-->	<ID, 'Mail'>

- ▶ Input for subsequent stages: **token stream**:

```
<SELECT> <ID, 'Mail'> <DOT> <ID, 'sender'> ... <EQ> <STRING, 'I love you'>
```



## 0.5.2 Identify Valid SQL Syntax

- ▶ Only a subset of token streams represent syntactically valid SQL queries;
  - Streams like ...  $\langle \text{NUM}, -42 \rangle \langle \text{SELECT} \rangle \langle \text{FROM} \rangle \langle \text{FROM} \rangle \dots$  don't make sense;

- ▶ We need rules that identify valid token (sub)sequences:

*“A condition is either a boolean or a quantifier or a condition in parentheses.”*

*“A quantifier either starts with  $\langle \text{FORALL} \rangle$  or  $\langle \text{EXISTS} \rangle$ , followed by a variable, followed by  $\langle \text{IN} \rangle \dots$ ”*

- ▶ Define a formal set of rules, a **grammar**, that describes valid token sequences (need **alternatives** and **sequences**):

$$\begin{array}{l} \textit{condition} \rightarrow \textit{boolean} \\ \quad \quad \quad | \textit{quantifier} \\ \quad \quad \quad | \langle () \textit{condition} () \rangle \end{array}$$
$$\begin{array}{l} \textit{quantifier} \rightarrow \langle \text{FORALL} \rangle \textit{id} \langle \text{IN} \rangle \textit{query} \langle : \rangle \textit{query} \\ \quad \quad \quad | \langle \text{EXISTS} \rangle \textit{id} \langle \text{IN} \rangle \textit{query} \langle : \rangle \textit{query} \end{array}$$

### 0.5.3 Resolve the Meaning of Variables and Identifiers

- ▶ SQL offers a great deal of freedom when it comes to the usage of variables or attribute names:

```
SELECT  Mail.sender, country
        FROM  Mail, DNS
        WHERE Mail.domain = DNS.domain
           AND subject = 'I love you'
```

- ▶ It will be our task to **make references explicit** (and complain in case of **ambiguity** etc.), and **come up with attribute names** used in the query result.
- ▶ This step saves the user from (a) being too explicit and (b) inventing variable and attribute names:

```
SELECT  x.sender AS sender, y.country AS country
        FROM  Mail AS x, DNS AS y
        WHERE x.domain = y.domain
           AND x.subject = 'I love you'
```

## 0.5.4 Check Types and Schemas

- ▶ Syntactical correctness does not imply we're dealing with a meaningful query:
  - ...WHERE  $42 < \text{'foobar'}$  ...
  - ...FORALL  $v$  IN VALUES 'foo', 'bar' :  $v + 1 < 3$  ...
- ▶ A query with **attribute references** like

```
SELECT   $x.a, x.b + 1$ 
FROM     $S$  AS  $x$ 
```

is acceptable only if the tuples in  $S$  are of type  $(\dots, a : \alpha, \dots, b : \mathbb{N}, \dots)$ .

- The query compiler will have to communicate with the DBMS's **schema manager** to check for schemas and attribute types.

## 0.5.5 How Could Query Evaluation Look Like?

```
SELECT  Mail.sender, country
FROM    Mail, DNS
WHERE   Mail.domain = DNS.domain
AND     subject = 'I love you'
```


- ▶ Something like the pseudo-code program below could evaluate the given query:

```
FOREACH  $y \in$  DNS DO
  FOREACH  $x \in$  Mail DO
    IF  $x.domain = y.domain \wedge x.subject = 'I love you'$  THEN
       $t \leftarrow (x.sender, y.country);$ 
      output  $t;$ 
```

- ▶ Execution is mainly driven by **Nested Loops** processing;
- ▶ this leads to a complexity of  $O(|Mail| \times |DNS|)$  ...  
... which becomes unacceptable for complex queries.

## 0.5.6 One Query, Many Programs

- ▶ The program below clearly is **more efficient** yet computes the **same query result**:


 Rewrite the Query Program.

- ▶ We have just used two major query optimization techniques:
  - Join Reordering, and
  - Selection Pushdown.


## 0.5.7 Query Operators

► Typical program patterns recur again and again.

■ Selection (with respect to predicate<sup>2</sup>  $p$ ):



■ Join (with respect to predicate  $p$ ):

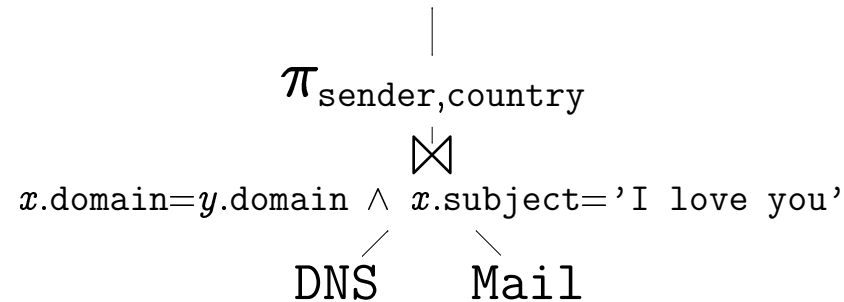


---

<sup>2</sup>A predicate is a function returning a result of type  $\mathbb{B}$ , e.g.,  $\leq :: (\mathbb{N}, \mathbb{N}) \rightarrow \mathbb{B}$ .

## 0.5.8 Query Operator Trees and Rewriting

- ▶ Query operators are assembled to form **query operator trees**:

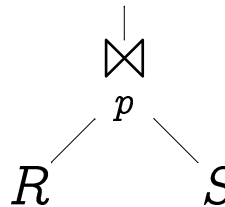
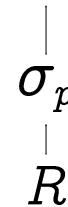


- ▶ With the help of operator trees we have a natural way of expressing **query rewriting steps**:



## 0.5.9 The Cheaper, the Better: Query Cost Models

- ▶ How can we guide the query optimizer in making sensible rewriting decisions?
  - Answer: **Rules of Thumb** and **Cost Models**.
- ▶ Assign a cost measure to each operator estimating its CPU and I/O activity:

	$cost(R \bowtie_p S) =$ $+  R $ $+  R  \times  S $ $+  R  \times  S  \times s_p$		$cost(\sigma_p(R)) =$ $+  R $ $+  R  \times s_p$
---	--	---	--

(with  $0 \leq s_p \leq 1$  denoting the selectivity of predicate  $p$ ).


### Cost Estimation is Sufficient

We do not care for the actual CPU cycles or number of I/O bursts as long as the cost estimation enables the optimizer to identify the cheaper of two operator trees.



## 0.5.10 Executing Operator Trees on a Machine

- ▶ Each operator implemented by separate subroutine in a query interpreter;
  - operators schedule each other via a simple *init–next–done?* interface:



- tuple-at-a-time scheme; avoids materialization of temporary results.
- ▶ Linearize operator tree in memory; nodes represented by heap addresses; operator invocation  $\equiv$  CPU jmp instruction.

To conclude this tour, this is a list of **high-level topics** that we will encounter during this course:

① **Syntactical and Semantical Analysis**

- ▶ lexing, parsing, parse tree construction
- ▶ type checking, variable scoping

② **Query Compilation**

- ▶ internal query representation, query calculus and algebra

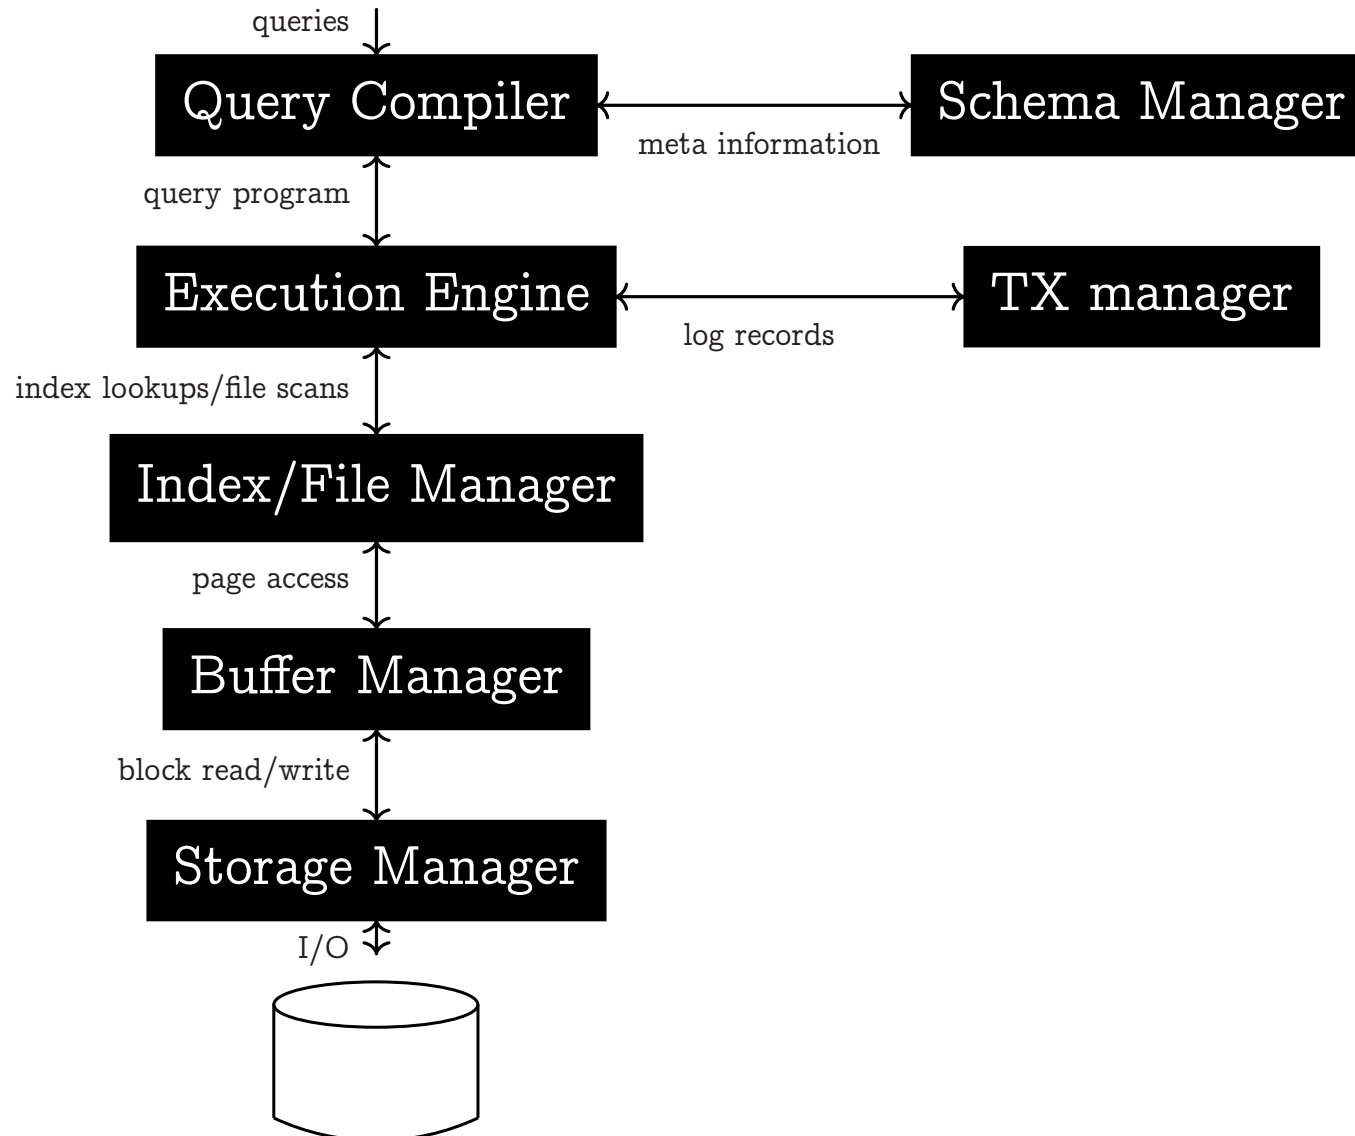
③ **Query Planning and Optimization**

- ▶ query rewriting, query costing, searching

④ **Query Execution**

- ▶ query interpreter design and implementation

# Components of a Typical DBMS



## 0.6 Recommended Reading

- [1] H. GARCIA-MOLINA, J. D. ULLMAN, AND J. WIDOM, *Database System Implementation*, Prentice Hall, New Jersey, 2000. ISBN 0-13-040264-8.
- [2] G. SAAKE AND A. HEUER, *Datenbanken: Implementierungstechniken*, MITP Verlag GmbH, 1999. ISBN 3-8266-0513-6.

# Contents

<b>0</b>	<b>Introduction to Query Compilation</b>	<b>1</b>
0.1	Welcome! . . . . .	1
0.2	Administrativa . . . . .	3
0.3	Some Remarks on these Slides . . . . .	4
0.4	Relational Databases and SQL . . . . .	6
0.5	A Guided Tour through an SQL Query Processor . . . . .	14
0.5.1	Character Streams and Tokens . . . . .	15
0.5.2	Identify Valid SQL Syntax . . . . .	16
0.5.3	Resolve the Meaning of Variables and Identifiers . . . . .	17
0.5.4	Check Types and Schemas . . . . .	18

0.5.5	How Could Query Evaluation Look Like? . . . . .	19
0.5.6	One Query, Many Programs . . . . .	20
0.5.7	Query Operators . . . . .	21
0.5.8	Query Operator Trees and Rewriting . . . . .	22
0.5.9	The Cheaper, the Better: Query Cost Models . . . . .	23
0.5.10	Executing Operator Trees on a Machine . . . . .	24
0.6	Recommended Reading . . . . .	27